

Enclosures:

1. Final Technical Report
2. "Sequential Decision Analysis for Non-Stationary Stochastic Processes"
by Brian Schaefer.

Copies sent to:

5 copies to NASA Scientific and Technical Information Facility
P. O. Box 33
College Park, Maryland 20740

2 copies to National Aeronautics and Space Administration
Proposal Control Officer
Office of University Affairs
Washington, D. C. 20546


2 copies to C. Ray Davis, Jr., Grants Officer
National Aeronautics and Space Administration
Langley Research Center
Hampton, Virginia 23365

1 copy to John H. Painter
National Aeronautics and Space Administration
Langley Research Center
Hampton, Virginia 23365

(NASA-CR-140625) SEQUENTIAL DECISION
ANALYSIS FOR NONSTATIONARY STOCHASTIC
PROCESSES Final Technical Report
(Northwestern Univ.) 16 p HC \$4.00

N74-34984

Unclas
CSCL 12A G3/19 51111



Final Technical Report

on NASA Grant

NGR 14-007-129

"Decision Directed Adaptive Estimation"

submitted by the Principal Investigator

Professor Brian Schaefer

Department of Industrial Engineering
and Management Sciences

Northwestern University

Evanston, Illinois 60201

October 1974

1. Progress of the Ph.D. Student Supported by this Grant.

Mr. Bombran Shetty, a Ph.D. student in the Industrial Engineering and Management Sciences Department at Northwestern University was supported (tuition and stipend) by this grant. He was unable to make significant progress toward the solution of the problem addressed in the research proposal, although he spent a great deal of time on it and worked very diligently. Finally, in March 1974, he was forced to drop out of school for personal reasons. This unfortunate series of circumstances caused considerable delay in the progress of the research.

2. Sequential Decision Analysis.

The principal investigator began to work full time on the grant after the departure of Mr. Shetty. It was felt that in order to develop an adaptive estimator for processes in which the mean and variance of the observation noise are unknown and may be changing in time, a procedure must be developed for making sequential decisions on non-stationary stochastic processes. Current statistical decision theory deals only with time independent random variables, and the results of optimal stochastic control theory, which do deal with the above problem, are usually not amenable to actual algorithmic implementation. Research toward development of such a procedure produced some independently interesting results, and are contained in the accompanying paper. These results also solve a major portion of the problem addressed in the research proposal. This paper is being submitted for publication to the Institute of Mathematical Statistics, and will be presented at the 1975 ORSA meeting in Chicago.

3. Continuing Work.

The principal investigator is continuing research on the problems addressed in the research proposal for this grant. One paper on model evaluation is currently being revised and the use of the above sequential decision algorithm in adaptive Kalman filtering is being considered. In the Kalman filter, the losses incurred for using an incorrect model are well-known, and these will be used as the loss function in the decision algorithm. Results of this research will be forwarded to NASA, as an addendum to this final report, when they are completed.

SEQUENTIAL DECISION ANALYSIS
FOR
NON-STATIONARY STOCHASTIC PROCESSES *

by

Brian Schaefer
Northwestern University
Evanston, Illinois

September 1974

Abstract

A formulation of the problem of making decisions concerning the state of non-stationary stochastic processes is given. An optimal decision rule, for the case in which the stochastic process is independent of the decisions made, is derived. It is shown that this rule is a generalization of the Bayesian likelihood ratio test; and an analog to Wald's sequential likelihood ratio test is given, in which the optimal thresholds may vary with time.

* Research supported in part by NASA Grant NGR 14-007-129.

I. Introduction

The general framework of the sequential decision problem has remained in essentially the same form as originally formulated by Wald [1947]. This formulation involved a sequential decision concerning the choice of obtaining another sample or making a final decision.

In this paper we generalize this problem to include making decisions on the state of a non-stationary stochastic process and are able to obtain a convenient solution for the case in which the state of the process is independent of the decisions made.

Such a formulation is of interest in problems involving estimation or signal detection as used in the tracking of missiles or commercial aircraft. In these problems, the decisions usually do not influence the original process. Another area in which this formulation is appropriate would be problems involving economic decisions where the processes under observation, such as stock prices or government indices, are relatively independent of decisions made on a personal or corporate level.

II. Sequential Decisions for Stochastic Processes

We will let T be a linearly ordered parameter set, and will assume that $z(t)$, $t \in T$ is a stochastic process defined on a probability space (Θ, β, γ) . We will also let $y(t)$, $t \in T$ be a stochastic process defined on a family of probability spaces $\{(\Omega, B, P_\theta): \theta \in \Theta\}$.

The set of admissible actions will be given as a measurable space (\mathcal{A}, A) . An action process, $a(t)$, $t \in T$, will be defined, on \mathcal{A} ; such a process is similar to a stochastic process without the probability measure, that is, $a: \mathcal{A} \rightarrow R^T$.

A measurable loss function $L: \Theta \times \mathcal{A} \rightarrow R$ will be defined as will a set, D , of measurable decision functions, where $d \in D$ and $d: \Omega \rightarrow \mathcal{A}$.

The function d will be required to operate on Ω only as a causal function of $y(t)$, $t \in T$; that is, for any $\tau \in T$

$$a(\tau) = d(y(t), t \leq \tau)$$

The Bayesian decision problem consists then in finding a $d^* \in D$ such that the risk function $r(d) = E[L(\theta, d(w))]$ is minimized

$$(1) \quad r(d^*) = \min_{d \in D} \int_{\Theta} \int_{\Omega} L(\theta, d(w)) dP_{\theta}(w) d\eta(\theta)$$

The complexity of the above minimization problem is determined by the nature of the probability spaces Θ and Ω , the loss function L , and the decision set D . The loss function L will be defined on Θ through the process $z(t)$, $t \in T$, so that $L(\theta, d(w)) = L(z(t), t \in T; d(y(t), t \in T))$.

If $z(t)$, $t \in T$, is a function of $d(y(t), t \in T)$ then the minimization in (1) is a problem usually studied in stochastic optimal control theory, see for example, Kushner [1967]. In the particular case when $z(\tau + \Delta)$ is a function only of $z(\tau)$ and $a(\tau)$ the problem is usually referred to as a Markov decision problem.

Although in the study of stochastic control theory and Markov decision processes it is possible to obtain necessary conditions for the optimal decision rule in some cases, these conditions often do not lead to a practical explicit solution. In the next section we will make several assumptions that will lead to a simple explicit solution to (1). These assumptions will usually be true in the case where the problem involves decisions concerning the state of the $z(t)$, $t \in T$ process, rather than the control of this process. In Wald's original formulation the process $z(t)$, $t \in T$ is a constant $z(t) = \alpha_1$, or α_2 for all $t \in T$, and the loss function is a constant until a decision is made, representing the cost of observations, and zero after the decision has been made. Although Wald did not adopt the Bayesian context, his results

would be unaltered if equal a priori probabilities were assumed. In the following section we will derive an interesting analog to Wald's results.

III. Non-Controlled Processes and Independent Action Processes

In this section we will assume that $z(t)$, $t \in T$ is not a function of $a(t)$, $t \in T$ and will consider the following form of the loss function

$$(2) \quad L(\theta, a) = \int_T L(z(t), a(t)) dt$$

where $L(z(t), a(t)) \geq 0 \quad \forall t \in T$

$$= \int_T L(z(t), d(y(\tau), \tau \leq t)) dt$$

The notation $f(z(t))$ will imply that $f(\cdot)$ may be a function of t as well as the value of z at t ; that is, $f(z(t)) = f(z(t), t)$. Assuming that the following integrals exist, the risk function, (1), becomes, with some abuse of notation,

$$\begin{aligned} r(d) &= \int_{\Theta} \int_{\Omega} \int_T L(z(t), d(y(\tau), \tau \leq t)) dt dP_{\theta}(w) d\eta(\theta) \\ &= \int_T \int_{\Omega} \int_{\Theta} L(z(t), d(y(\tau), \tau \leq t)) d\eta(\theta|w) dP(w) dt \\ &= \int_T \int_{R^t} \int_R L(z(t), d(y(\tau), \tau \leq t)) d\eta(z(t)|y(\tau), \tau \leq t) dP(y(\tau), \tau \leq t) dt \end{aligned}$$

We will let $\mathcal{A}(t)$ represent the set of admissible actions, at time t , and we will make the assumption that $\mathcal{A}(t)$ is independent of $a(\tau)$, $t \neq \tau$ for any $t, \tau \in T$. That is, the set of admissible actions at any given time does not depend on an action taken at any other time. Such processes $a(t)$, $t \in T$, we will call independent action processes

If we let

$$(3) \quad \mathcal{L}(\alpha, y, t) = \int_R L(z(t), d(y(\tau), \tau \leq t) = \alpha) d\eta(z(t)|y(\tau), \tau \leq t), \alpha \in \mathcal{A}(t)$$

then the above risk function becomes

$$(4) \quad r(d) = \int_T \int_R \mathcal{L}(\alpha, y, t) dP(y(\tau), \tau \leq t) dt$$

Theorem

If: $a(t)$, $t \in T$ is an independent action process; $z(t)$, $t \in T$ does not depend on $a(t)$, $t \in T$; and the loss function is of the form given in (2), then the risk (1) is minimized by the following decision rule:

$$d(y(\tau), \tau \leq t) = \alpha^* \quad \alpha^* \in \mathcal{A}(t)$$

iff

$$(5) \quad \mathcal{L}(\alpha^*, y, t) \leq \mathcal{L}(\alpha, y, t) \quad \forall \alpha \in \mathcal{A}(t)$$

Proof

Since $L(z(t), d(y(\tau), \tau \leq t)) \geq 0 \quad \forall t \in T$, from (3) we have $\mathcal{L}(\alpha, y, t) \geq 0$ for all $\alpha \in \mathcal{A}(t)$, $t \in T$, and all $y(\tau)$, $\tau \leq t$. Thus (4) will be minimized by choosing the $\alpha \in \mathcal{A}(t)$ that minimizes $\mathcal{L}(\alpha, y, t)$ for each $t \in T$ and each $y(\tau)$, $\tau \leq t$. Given (4), this last statement, may be proved simply by contradiction. □

We will also define

$$(6) \quad Q(\alpha, y, t) = \int_R L(z(t), d(y(\tau), \tau \leq t) P(y(\tau), \tau \leq t | z(t)) d\eta(z(t))$$

then

Corollary

Given the assumptions of the above theorem, the risk (1) is minimized by the decision rule

$$d(y(\tau), \tau \leq t) = \alpha^* \quad \alpha^* \in \mathcal{A}(t)$$

iff

$$(7) \quad Q(\alpha^*, y, t) \leq Q(\alpha, y, t) \quad \forall \alpha \in \mathcal{A}^t$$

Proof

From (3) and (6) $Q(\alpha, y, t) = \mathcal{L}(\alpha, y, t)P(y(\tau), \tau \leq t)$

and therefore

$$Q(\alpha^*, y, t) \leq Q(\alpha, y, t)$$

iff

$$\mathcal{L}(\alpha^*, y, t) \leq \mathcal{L}(\alpha, y, t)$$



a) Estimation Given a Quadratic Loss Function

If

$$\begin{aligned} L(z(t), d(y(\tau), \tau \leq t)) \\ = (z(t) - d(y(\tau), \tau \leq t))^2 \end{aligned}$$

then the optimal decision rule as determined from (3) and (5) is

$$(8) \quad d(y(\tau), \tau \leq t) = E[z(t) | y(\tau), \tau \leq t].$$

This result is given in Doob [1953].

The conditions required of $y(t)$ and $z(t)$, $t \in T$ by the Kalman filter, Kalman [1960], are precisely those such that (8) can be computed recursively in time.

b) Finite State and Action Spaces

Suppose

$$\begin{aligned} z(t) &\in \{\alpha_i : i = 1, \dots, n\} & \forall t \in T \\ \text{and} \quad a(t) &\in \{\beta_i : i = 1, \dots, m\} & \forall t \in T \end{aligned}$$

and that the observation process $y(t)$, $t \in T$ is some process related to $z(t)$, $t \in T$, as depicted in Figure 1.

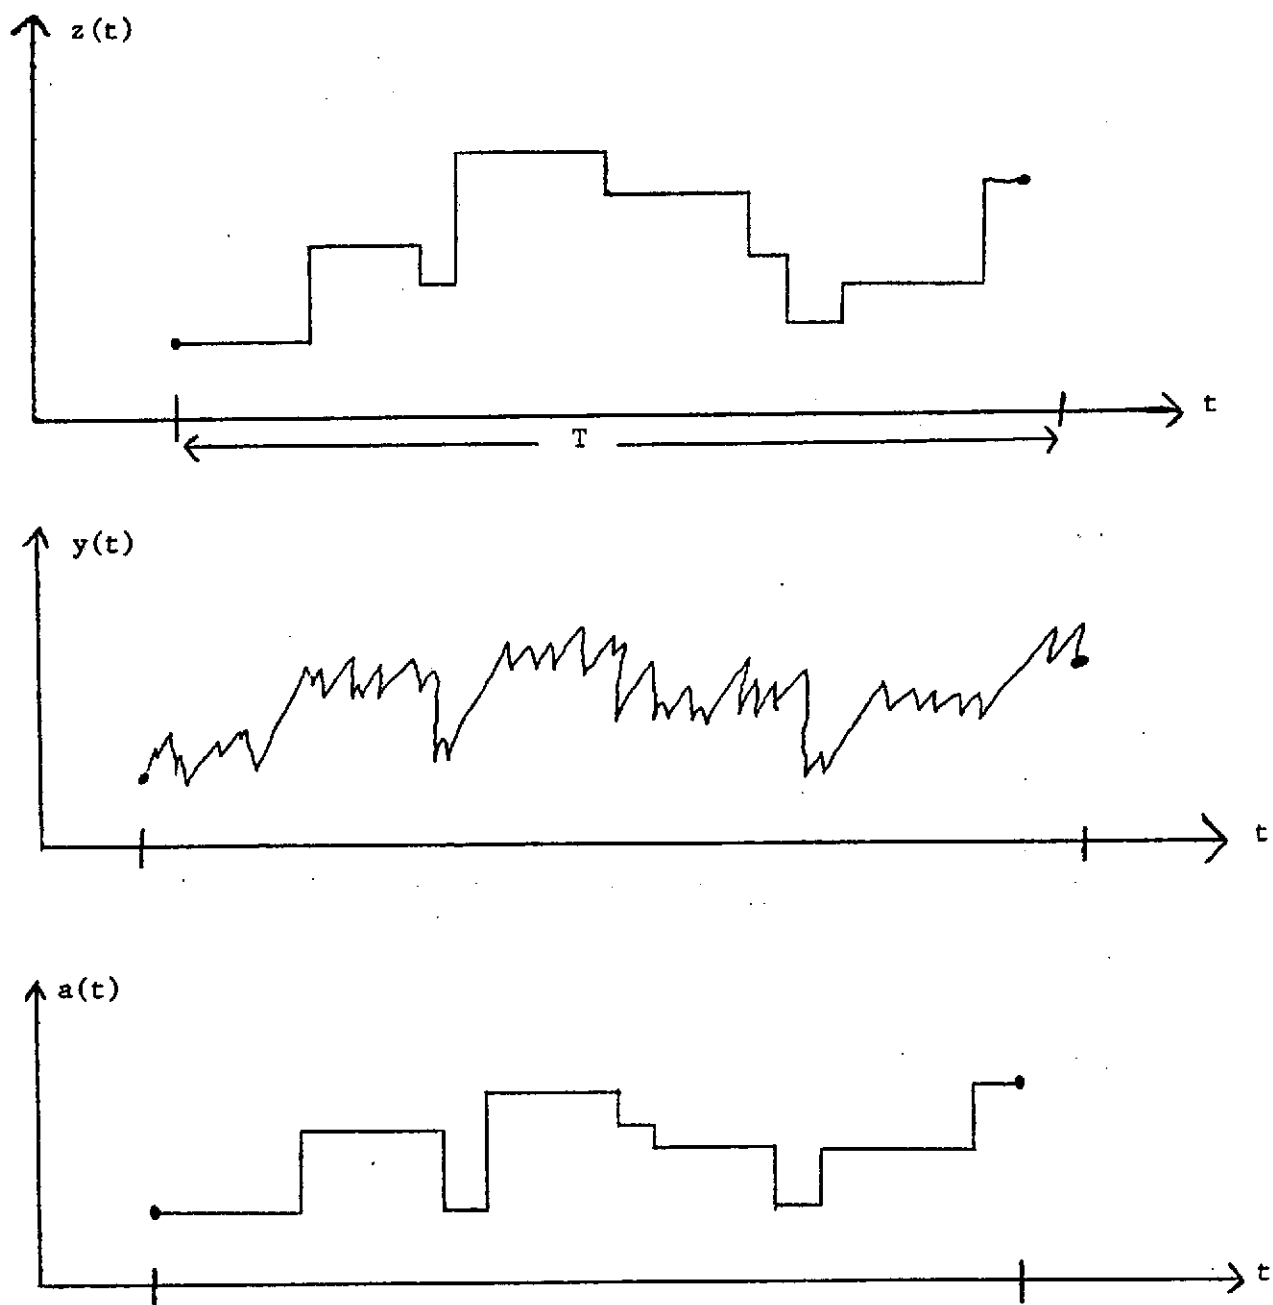


Figure 1: Sample realizations of $z(t)$, $y(t)$, $a(t)$, $t \in T$

$a(t)$ is a function of $y(\tau)$, $\tau \leq t$, and we will define

$$L(z(t) = \alpha_i, d(y(\tau), \tau \leq t) = \beta_j) = L_{ij}(t) \geq 0.$$

In a practical situation we might be trying to determine the state of $z(t)$ from noisy data, and the loss would be minimized if $a(t) = z(t)$, $t \in T$.

The above Theorem gives the following optimal decision rule.

$$a(t) = \beta_{j*}$$

$$\text{iff } \sum_{i=1}^n L_{ij*}(t) \eta(z(t) = \alpha_i | y(\tau), \tau \leq t) \leq \sum_{i=1}^n L_{ij}(t) \eta(z(t) = \alpha_i | y(\tau), \tau \leq t)$$

$$j = 1, \dots, m.$$

(i) Likelihood Ratio Test

It is well known that for simple binary random variables and fixed sample size, both the optimal Bayes test and the Neyman-Pearson test result in comparing the likelihood ratio with a simple threshold. If $z(t) \in \{\alpha_1, \alpha_2\}$, and $a(t) \in \{\alpha_1, \alpha_2\} \forall t \in T$ and assuming that $L_{ii}(t) \leq L_{ij}(t)$, $i = 1, 2$, then from (7) the optimal decision rule is shown to be

$$a(t) = \alpha_1 \quad \text{iff} \quad \frac{P(y(\tau), \tau \leq t | z(t) = \alpha_1)}{P(y(\tau), \tau \leq t | z(t) = \alpha_2)} \geq \frac{[L_{21}(t) - L_{22}(t)]\eta(z(t) = \alpha_2)}{[L_{12}(t) - L_{11}(t)]\eta(z(t) = \alpha_1)}$$

$$a(t) = \alpha_2 \quad \text{otherwise.}$$

Thus if $z(t)$ is constant in time, and the loss function is independent of time, the above decision rule reduces to the familiar likelihood ratio test. On the other hand, the result shows that for a stochastic process the optimal decision rule consists of a time varying likelihood ratio test.

(ii) Analog to Wald's Original Formulation

In Wald's original formulation, he considered processes which were in one of two states for all $t \in T$; that is, the process was a simple

binary random variable. This will be a special case of the following, in which we will assume that the process may be in either of two states at any time $t \in T$.

$$z(t) \in \{\alpha_1, \alpha_2\} \quad \forall t \in T$$

We will let our action space be such that

$$a(t) \in \{\alpha_1, \alpha_2, \alpha_3\} \quad \forall t \in T$$

where:

$a(t) = \alpha_1$ corresponds to the decision that $z(t) = \alpha_1$;

$a(t) = \alpha_2$ corresponds to the decision that $z(t) = \alpha_2$;

and, $a(t) = \alpha_3$ corresponds to the action of making no decision at this time, other than to wait until we have obtained another observation. We will specify a loss function that requires us to pay for this additional information and waiting time. We will let

$$L_{11}(t) < L_{13}(t) < L_{12}(t) \quad \forall t \in T$$

and

$$L_{22}(t) < L_{23}(t) < L_{21}(t) \quad \forall t \in T.$$

That is, the loss incurred for making a correct decision is less than the loss for making no decision, which is less than the loss for making the wrong decision. For simplicity we will normalize the losses such that

$$L_{11}(t) = L_{22}(t) = 0.$$

If we assume that $P(z(t) = \alpha_1) = P(z(t) = \alpha_2) \quad \forall t \in T$ then the decision rule determined by (7) consists of computing the likelihood ratio

$$(9) \quad \Lambda(y, t) = \frac{\eta(y(\tau), \tau \leq t | z(t) = \alpha_1)}{\eta(y(\tau), \tau \leq t | z(t) = \alpha_2)}$$

and setting

$$a(t) = \alpha_1 \text{ if } \Lambda(y, t) > \frac{L_{21}(t) - L_{23}(t)}{L_{13}(t)}$$

$$a(t) = \alpha_2 \text{ if } \Lambda(y,t) < \frac{L_{23}(t)}{L_{12}(t) - L_{13}(t)}$$

$$a(t) = \alpha_3 \text{ otherwise.}$$

If the loss function is time independent, then the similarity of this to Wald's result of constant thresholds is striking. In particular if we let $L_{12} = L_{21}$ and $\frac{L_{23}}{L_{21}} = \text{Probability of a Type I error.}$

$$\frac{L_{13}}{L_{21}} = \text{Probability of a Type II error.}$$

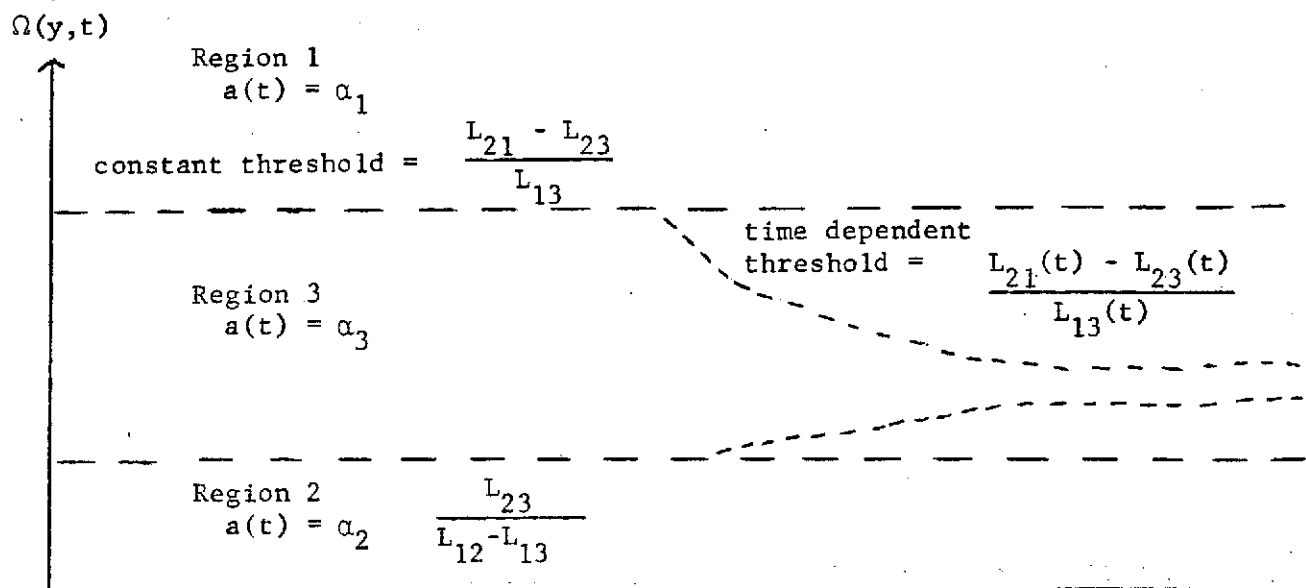
then these are exactly the Wald thresholds.

If we wish, however, we can make $L_{13}(t)$ and $L_{23}(t)$, increasing functions of time, which correspond to making the loss incurred by indecision greater the longer a decision is delayed. This would bring the above thresholds closer together as shown in Figure 2, although we note that the inequality

$$L_{12}(t) > L_{13}(t) + L_{23}(t) \quad \forall t \in T$$

must be satisfied.

This inequality is a statement of the fact that if the total cost of indecision is greater than the cost of a wrong decision, then a decision should always be made.



A major difference between this test and Wald's test is, that the test continues for a fixed time interval, T . Once a threshold is crossed, the test does not stop, rather the action is constant until the same threshold is recrossed in the opposite direction.

IV. Conclusion

This paper has considered the problem of making decisions on the state of a stochastic process. A solution to the problem has been derived from the case in which the state of the process is independent of the decisions made, the set of admissible actions at any time is independent of the action taken at any other time, and the loss function is of the form given in (2). These assumptions are usually made implicitly in the derivation of the conditional expectation as the solution to the minimum variance estimator, and this solution is shown to follow from the decision rule derived in this paper. The above assumptions are also true in the formulation of the standard Bayesian likelihood ratio test, and the Neyman-Pearson test, and this paper therefore, is a generalization of these tests.

In the Wald test, the assumption of the set of admissible actions being independent of actions taken at any other time, is not true, since once a threshold is crossed no more observations may be taken. Often, however, one would like to formulate a statistical test, in a sequential manner, and continue to accept observation after one hypothesis has tentatively been accepted. This would be particularly true if the hypothesis that was true, could change over time. The solution to such a formulation is given in this paper.

References

- Doob, J. L. (1953), Stochastic Processes, Wiley, New York.
- Kalman, R. (1960), A new approach to linear filtering and prediction,
ASME, J. Basic Engr., D82, 35-44.
- Kushner, H. J. (1967), Stochastic Stability and Control, Academic Press,
New York.
- Wald, A. (1947), Sequential Analysis, Wiley, New York.